RECEIVED
MAY 31 1973

Postel

NIC 16443
NWG/RFC 512

Wayne Hathaway (AMES-67)
25 MAY 1973

MORE ON LOST MESSAGE DETECTION

I would like to second Edwin Meyer's (RFC#492) strong opposition to the proposals made in RFC#467 concerning solutions to the "lost allocate" and "half-closed" phenomena. In particular I support all of his principles concerning the "half-closed" phenomenon. I also agree that the proposed "lost allocate" solution tends to mask the real problem of lost messages. I would, however, like to propose the following alternative scheme for recognizing lost messages.

I propose that one of the two unused eight-bit bytes in the level 2 message leader be designated the "sequence Control Byte" (SCB). This SCB would be essentially a modulo 255 message count. Upon receipt of a message, the receiving NCP would compare the SCB in the message with the expected SCB as computed from the SCB in the previous message on the same link. A discrepancy indicates a lost message, which could then be reported immediately via an appropriate ERR message. This ERR message (to be defined) would contain both received and expected SCB's, allowing possible recovery of the lost message (if sufficient space were available in the sending host to save the last several messages for each link). At any rate, the lost message would be recognized immediately, whether it was an ALL (or any control message) or a data message. The message with the unexpected SCB should be processed normally, with the SCB for the next message being computed from it.

For compatibility, the SCB would be defined such that an SCB of zero indicates that no checking is to be done. The SCB following 255 would thus be 1. This would mean that current NCP's would not have to be changed unless actual checking were desired (since the level 2 protocol specifies that these two unused bytes must be zero). This special definition of a zero SCB would also allow RST's and ERR's to bypass checking, which would be useful in avoiding possible loops.

This proposed scheme is similar to the second scheme suggested by Jon Postel (RFC#516) except that it is on a per-link basis rather than a per-host basis. This is significant, however, as it removes the requirement that all messages from one host to another arrive in the order sent (which cannot be guaranteed). It also provides for compatibility with existing NCP's. Jon's first proposal (save all messages until RFNM received) is weak in two areas: first, it is possible that the receiving IMP has sent a RFNM for a message that in fact never gets to its host, and second, it requires (at least for swapped systems such as ours) either that messages be saved in resident storage (expensive) or that RFNM's be handled by a swapped process (also expensive). The third proposal (that of a host-to-host acknowledgment scheme) is perhaps the best, but as that requires quite major changes to the level 2 protocol, an interim solution such as that proposed here seems of value.